

What is claimed is:

1. A method for storing data on an array of disk drives, comprising:
 - dividing extents of data into data segments;
 - defining a parity group comprising at least one data segment from a plurality of data segments from different extents;
 - generating parity segments from said data segments in said parity group;
 - storing said parity segments;
 - generating a second parity segment from said parity segments; and
 - storing said second parity segment.
2. The method of claim 1 wherein said storing said parity segments step further comprises the step of appending said parity segments after said data segments in each extent.
3. The method of claim 2 wherein each of n-1 extents in said parity group has one parity segment appended thereon.
4. The method of claim 1 wherein said storing said second parity segments step further comprises the step of appending said second parity segment to one of n extents in the parity group.
5. The method of claim 4 wherein said storing said parity segments step further comprises the step of appending said second parity segment after said data segments in the nth extent.
6. The method of claim 1 wherein said generating said parity segments further comprises the step of storing data information from one segment in each extent of said parity group.
7. The method of claim 6 wherein said generating step further comprises the step of storing the data information sequentially from each data segment in each extent of said parity group.

8. The method of claim 1 wherein said generating said parity segments further comprises the step of storing parity information from one parity segment in each extent of said parity group.

9. The method of claim 8 wherein said generating step further comprises the step of storing the parity information sequentially from each parity segment in each extent of said parity group.

10. A method for reconstructing lost data from a failed disk drive in an array of disk drives, said disk drive array formed into at least one parity group, wherein each parity group comprises data extents divided into data segments and parity segments that are generated from the data segments, said method comprising:

- retrieving the parity segments from operative disk drives in a parity group containing the failed disk drive;
- reconstructing parity information for the failed disk drive from the parity segments on the operative disk drives;
- reconstructing data segments for the failed disk drive from the reconstructed parity information, the parity segments, and data segments on the operative disk drives.

11. The method of claim 10 wherein said retrieving step further comprises storing all of the data segments and the parity segments of the operative disk drives of the parity group in a buffer memory.

12. An apparatus for storing data comprising:

- a plurality of disk drives formed into a parity group having m -data extents, wherein each disk drive comprises $1/m$ data extents further divided into p -data segments;
- $m-1$ parity segments respectively coupled to $m-1$ data extents; each $m-1$ parity segment storing parity information corresponding to m -data segments sequentially positioned across the m -data extents located on different disk drives in the parity group; and

a second parity segment containing parity information for the $m-1$ parity segments and coupled to one of the m -data extents.

13. The apparatus of claim 12 wherein said parity group further comprises $m-p$ - data segments.

14. The apparatus of claim 12 wherein said $m-1$ parity segments are respectively appended at the end of said $m-1$ data extents.

15. The apparatus of claim 14 wherein said $m-1$ parity segments are sequentially appended to the m -data extents across the disk drives in the parity group.

16. The apparatus of claim 12 wherein said second parity segment is appended at the end of one of said $m-1$ data extents.

17. The apparatus of claim 16 wherein said second parity segment is appended to the m^{th} data extent.

18. An apparatus for storing data on an array of disk drives, comprising:
means for defining at least one parity group from said array of disk drives;

means for dividing extents of data from said at least one parity group into data segments;

means for generating parity segments from said data segments in said parity group;

means for generating a second parity segment from said parity segments, wherein said parity segments from said data segments and said second parity segment are attached to different data extents in the parity group.

19. The apparatus of claim 18 wherein the means for defining at least one parity group further comprises apportioning said array of disk drives into equally sized parity groups each having an equal number of disk drives.

20. The apparatus of claim 18 wherein the means for dividing extents of data further comprises defining one extent per disk drive in the parity group.

21. The apparatus of claim 19 wherein the data segments are apportioned equally in each extent of the at least one parity group.

22. The apparatus of claim 18 wherein said parity segments and said second parity segment are attached at the end of each extent.

23. A method for recovering data from a failed disk drive in a disk array comprising a plurality of disk drives, said plurality of disk drives having a plurality of extents apportioned equally into a plurality of parity groups, said method comprising:

- striping data from at least one file sequentially across the plurality of extents of the plurality of parity groups;
- identifying a failed parity group containing the failed disk drive;
- outputting, in realtime to users in a normal disk access mode, said data from parity groups without the failed disk drive;
- reconstructing, in a parity correction mode of operation, said data from the at least one failed parity group; and
- outputting, in realtime to the users in the failed parity group, said reconstructed data.

24. The method of claim 23 wherein said striping step further comprises:

- dividing each extent in each parity group into a plurality of data segments;
- storing data information in the plurality of data segments; and
- forming a plurality of parity segments from the data segments in the parity group; and
- appending one parity segment to the end of each extent, wherein parity segments in the parity group respectively store parity information corresponding to the parity group in which the parity segments reside.

25. The method of claim 24 wherein the parity segments in the parity group store parity information for data segments in the parity group, as read sequentially across each extent, data segment by data segment.

26. The method of claim 25 wherein one of the parity segments in the parity group is a pure parity segment, which only contains parity information for other parity segments in the parity group.

27. The method of claim 26 wherein the pure parity segment in the parity group is positioned at the end of the last extent in the parity group.

28. The method of claim 7 wherein user access to the at least one file comprises accessing each extent sequentially, extent by extent across the plurality of disks in the disk array.

29. The method of claim 28 further comprising the step of accessing only the data segments of each extent in the parity group during the normal disk access mode.

30. The method of claim 29 further comprising the step of accessing only the parity segments in the parity group during the parity correction mode.

31. The method of claim 30 wherein the identifying step further comprises: monitoring performance of the plurality of disks drives in the disk array; comparing performance of each disk drive to a minimum threshold; declaring a disk drive performing below said minimum threshold as being the failed disk; and stopping access to the failed disk drive.

32. The method of claim 31 wherein, in an instance where the failed disk has been identified, the method further comprises transitioning the users accessing the failed parity group in the normal disk access mode to the parity correction mode.

33. The method of claim 32 wherein the parity correction mode comprises the step of implementing a disk regeneration algorithm to recover extent data on the failed disk drive from extent data on operable disks in the parity group.

34. The method of claim 33 wherein the disk regeneration algorithm comprises the steps of:

retrieving the data segments on the operable disk drives within the failed parity group;

regenerating failed data segments from the data segments on the operable disk drives and the parity segments in the parity group; and

buffering the regenerated extent data and the extent data from the operable disk drives in the parity group in a buffer memory.

35. The method of claim 34 wherein the user accesses the buffer memory containing the extent data from the parity group during the parity correction mode.

36. The method of claim 35 wherein the user is transitioned back to the normal disk access mode for accessing extent data in the parity groups without the failed disk drive.

37. The method of claim 36 wherein during the normal disk access mode the user sequentially accesses the extent data in the disk drives of the parity groups without the failed disk drive.

38. The method of claim 37 further comprising admitting new users access to the plurality of disk drives in the disk array coupled to the server according to a disk access admission algorithm having stochastic parameters, which include accounting for accessing the parity segments in the failed parity group during the parity correction mode.

39. The method of claim 38 wherein accounting for accessing the parity segments in the failed parity group further comprises setting expected disk access times greater than actual disk access times.